THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Indexability and index heuristics for a simple class of inventory routing problems

Citation for published version:

Black, D, Archibald, T & Glazebrook, K 2009, 'Indexability and index heuristics for a simple class of inventory routing problems', *Operations Research*, vol. 57, no. 2, pp. 314-326.
<https://doi.org/10.1287/opre.1070.0505>

Digital Object Identifier (DOI):

[10.1287/opre.1070.0505](https://doi.org/10.1287/opre.1070.0505)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Operations Research

Publisher Rights Statement:

© INFORMS 2009. Permission has been granted by INFORMS to publish the full-text publisher's PDF of this article on the University of Edinburgh Research Explorer; 16/07/2013. Originally published on INFORMS Online: <http://or.journal.informs.org/content/57/2/314.abstract>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Indexability and Index Heuristics for a Simple Class of Inventory Routing Problems

T. W. Archibald

Management School, University of Edinburgh, Edinburgh, EH8 9JY, United Kingdom, t.archibald@ed.ac.uk

D. P. Black, K. D. Glazebrook

Management School, Lancaster University, Lancaster, LA1 4YX, United Kingdom
{dan.black@lancaster.ac.uk, k.glazebrook@lancaster.ac.uk}

We utilise and develop Whittle's restless bandit formulation to analyse a simple class of inventory routing problems with direct deliveries. These routing problems arise from the practice of vendor-managed inventory replenishment and concern the optimal replenishment of a collection of inventory holding locations controlled centrally by a decision maker who is able to monitor inventory levels throughout the network. We develop a notion of location indexability from a Lagrangian relaxation of the problem and show that (subject to mild conditions) the locations are indeed indexable. We thus have a collection of location indices in closed form, namely, real-valued functions of the inventory level (one for each location), which measure in a natural way (namely, as a fair charge for replenishment) each location's priority for inclusion in each day's deliveries. We discuss how to use such location indices to construct heuristics for replenishment and assess a greedy index heuristic in a numerical study where it performs strongly. A simpler approximate index analysis is available for the case in which the demand at each location is Poisson. This analysis permits a more explicit characterisation of the range of holding cost rates for which (approximate) location indexability is guaranteed.

Subject classifications: dynamic programming/optimal control: applications, models; inventory/production: policies; probability: stochastic model applications.

Area of review: Stochastic Models.

History: Received March 2007; revision received July 2007; accepted September 2007. Published online in *Articles in Advance* January 5, 2009.

1. Introduction

The classical work of Gittins (1979, 1989) on multiarmed bandit problems gave formal mathematical expression to the natural idea that many complex dynamic stochastic optimization problems have good (sometimes optimal) solutions that use state-based calibrations (*indices*) of the options facing the decision maker to assist in choosing between them. A recent survey describing a range of developments of Gittins' work is due to Mahajan and Teneketzis (2008).

One particularly important extension of Gittins' model was described by Whittle (1988). Whittle's restless bandit problems concern the development of rules for the optimal activation of collections of stochastic projects or bandits. In Whittle's model, M projects from a collection of size N ($>M$) are chosen for activation at each of a sequence of decision epochs. Projects may change state whether active or passive, though according to different dynamics. The innovation in Whittle's model was precisely the state evolution of passive projects, a phenomenon he called *restlessness*. Each project earns a state-dependent reward at each epoch. The goal of analysis is the determination of a policy or rule for project activation to maximise the average reward rate earned from all projects over an infinite horizon. Whittle's model is almost

certainly intractable, having been shown to be PSPACE-hard by Papadimitriou and Tsitsiklis (1999). Whittle used a Lagrangian relaxation of his restless bandit problem to develop an index heuristic for the subclass of problems in which each constituent project passes an indexability test. Whittle's heuristic attaches an index to each project, namely, a real-valued function on that project's state space, and chooses for activation at each decision epoch the M projects with the largest index values. Ties are broken arbitrarily. Weber and Weiss (1990, 1991) established a form of asymptotic optimality for Whittle's heuristic under given conditions, while Niño-Mora (2001, 2002) and Glazebrook et al. (2002) have explored the issue of indexability. Further, Whittle's approach has been shown to give rise to strongly performing policies in a range of application domains. These include the control of service systems, investment problems, and the dynamic routing of jobs for service and machine maintenance. See, for example, Ansell et al. (2003), Glazebrook et al. (2005, 2006), and Glazebrook and Kirkbride (2007).

This current paper is, to the authors' knowledge, the first to apply index theory to problems concerning inventory routing. In §2, we describe an inventory routing problem (IRP) which arises from the developing practice of vendor-managed inventory replenishment. The latter term

describes situations in which a central controller monitors inventory levels at several locations and makes decisions regarding inventory replenishment at those locations in the interests of the network as a whole. Surveys of research contributions to such problems can be found in Federgruen and Simchi-Levi (1995) and Kleywegt et al. (2002). The latter paper explains why the inventory routing problem with direct deliveries (IRPDD) in which the delivery trucks under the decision maker's control visit just a single location on each trip from a central depot is important. In §2, we describe a version of IRPDD in which, whenever a decision is made to replenish a particular location, it is replenished in full. Such a model has previously been considered by Barnes-Schuster and Bassok (1997). IRPs are extremely challenging and the solution methods proposed have often been complex and computationally demanding. In sharp contrast, this paper offers something simple and natural (albeit within a restricted class of models), namely, inventory control policies which are constructed by attaching an index to each location, a function of its current inventory level, with large indices indicating that the location concerned is a high priority for replenishment.

In §2, we develop and adapt Whittle's general approach to our IRP. We describe what it means that a location be indexable and, if so, how its index is defined. Some proposals are advanced concerning how index values can be used to construct heuristic policies for inventory replenishment. In §3, we show that in our problem, locations are guaranteed indexable when holding costs are ignored. We are able to infer that they remain indexable when holding cost rates are set at realistic levels. Location indices are given in closed form. For the important case of Poisson demands, we develop a simpler, but approximating, index analysis in §4. This approach allows us to formally specify a range of holding cost rates for which (approximate) indexability is achieved. The paper concludes in §5 with a numerical study in which index-based heuristic policies perform strongly.

2. The Model and Methodology

In our IRPDD, each of L locations holds supplies of an item. The inventory level of the item at each location is recorded at regularly spaced time epochs. These epochs will be referred to as "the beginning of each day" throughout the paper, although nowhere do we require the day to be our basic time unit. Once all L inventories have been inspected, a decision is made concerning which locations (if any) should be resupplied that day. The goal of decision making is to minimise a combination of inventory and delivery costs for the network. Further details are as follows:

(i) The daily demand for the item at location l is described by the probability mass function $\{p_{jl}, j \geq 0\}$, with p_{jl} the probability that j items will be requested at location l in a single day. We use

$$\lambda_l = \sum_{j=0}^{\infty} j p_{jl}$$

for the mean of this distribution and refer to it as the *demand rate* at location l . Daily demands are assumed independent across locations and over time.

(ii) We assume that *holding costs* and *shortage costs* are incurred at each location with h_l the holding cost rate (per item per day) and σ_l the cost incurred whenever a demand cannot be met through a shortage at location l . There is no backlogging of demand.

(iii) During each day, M identical trucks are available to make deliveries. Each truck makes a series of round trips between the central depot and individual locations. Resupplying location l incurs a fixed cost K_l together with a cost of C_l per item supplied, takes total delivery time τ_l (an integer multiple of the time taken for a single round trip to l) and results in S_l (the replenishment level) items being available to meet demand at location l the following day. We are here making simplifying modelling assumptions which cannot be met exactly in practice because of uncertainty in both the inventory level at a location at each delivery time and in the demand arising at a location in that part of the day which follows completion of each delivery. This could be overcome by a modelling approach which incorporated information on inter alia the timing of deliveries. Such an approach would involve major additional complexities and would, in our judgement, add little to the analysis. We also assume that delivered items are not available to meet demand at a location l until the beginning of the next day. Put briefly, we treat delivered items as though they arrived at the end of their day of delivery.

(iv) A set of locations will be said to be *feasible* if there exists a schedule of round trips to those locations which can be made by the trucks during a single day and which would guarantee a completed delivery to each. To identify a feasible set, we need to know, for each l , the time taken for a single round trip between the depot and location l (t_l) and hence the number of visits to l to complete a delivery (τ_l/t_l). Expressed differently, a set of locations is feasible if there exists an M -fold partition of the visits required to complete a delivery to each such that the visits in each partition set can be completed in one day by a single truck. We assume that each individual location constitutes a feasible set.

Please note that, throughout our numerical investigations, we take the cost parameters h_l , σ_l , and C_l to be l -independent. However, this is not required for the theory.

We write π for a general *delivery policy*, namely, a rule for choosing a feasible set of locations to be supplied each day. Our goal is to choose π to minimise an aggregate cost rate

$$C^{\text{opt}} = \inf_{\pi} \left\{ \sum_{l=1}^L C_l(\pi) \right\}, \quad (1)$$

where in (1), $C_l(\pi)$ is the average cost rate (including the cost of deliveries to l and inventory costs) incurred at l over an infinite horizon. By the theory of stochastic

dynamic programming (DP) (see, for example, Puterman 1994) we may restrict attention to the class of stationary delivery policies which make decisions on the basis of current inventory levels only. That said, direct application of DP for problems of realistic size (in particular, reasonable values of L) is not computationally possible. Our search is, therefore, for effective heuristic approaches.

We proceed in three steps, each of which involves the development of a (further) relaxation of the problem. In Step 1, we relax the problem by identifying actions (choices of locations for replenishment) with subsets of the set of locations $\{1, 2, \dots, L\}$ and by declaring subset $\mathcal{L}$ to constitute a feasible action if

$$\sum_{l \in \mathcal{L}} \tau_l \leq M. \quad (2)$$

Hence, in this relaxation we regard the M trucks as a single resource and dispense with the constraints imposed by the discreteness of each one. In Step 2, we further relax the problem by requiring that the resource constraints (2) only be satisfied over the infinite horizon in a time average sense. Hence, policy π will now be declared feasible if

$$\sum_{l=1}^L \tau_l I_l(\pi) \leq M, \quad (3)$$

where in (3), $I_l(\pi)$ is the rate at which deliveries are made to location l under stationary policy π over an infinite horizon. We write $\tilde{\Pi}$ for the set of stationary delivery policies satisfying (3) and

$$\tilde{C}^{\text{opt}} = \inf_{\pi \in \tilde{\Pi}} \left\{ \sum_{l=1}^L C_l(\pi) \right\} \quad (4)$$

for the cost rate minimised over $\tilde{\Pi}$. We now proceed to Step 3 and develop a Lagrangian relaxation of the optimization problem in (4) by incorporating terms in the objective which penalise violations of the (time average versions of) the resource constraints as expressed by (3). Hence, we now write

$$\tilde{C}^{\text{opt}}(\nu) = \inf_{\pi} \left[\sum_{l=1}^L \{C_l(\pi) + \nu \tau_l I_l(\pi)\} - \nu M \right]. \quad (5)$$

In (5), the infimum is over a class of policies which, at each decision epoch, can choose to replenish any number of locations with no resource constraint imposed. It is plain that $\tilde{C}^{\text{opt}}(\nu) \leq \tilde{C}^{\text{opt}} \leq C^{\text{opt}}$ for any $\nu \in \mathbb{R}^+$.

However, by the nature of both the class of policies over which the minimisation in (5) is taken and the separably additive nature of the objective, the Lagrangian relaxation in (5) admits an additive decomposition in which the optimization for the L -location problem is replaced by L single-location optimization problems. We write

$$\tilde{C}^{\text{opt}}(\nu) = \sum_{l=1}^L C_l^{\text{opt}}(\nu) - \nu M. \quad (6)$$

In (6), $C_l^{\text{opt}}(\nu)$ is the value of an optimization problem involving location l alone. In this problem, at the start of each day the inventory level at location l is observed and a decision is made concerning whether to resupply it or not. In making such decisions, the goal is to minimise a cost rate which is an aggregate of inventory/delivery costs for l ($C_l(\pi)$) and charges imposed for the resource consumed in making deliveries ($\nu \tau_l I_l(\pi)$). Here the Lagrange multiplier ν has an economic interpretation as a charge levied per unit of resource (delivery time) consumed. An optimal policy for the Lagrangian relaxation in (5) will simply run optimal policies for the individual locations in parallel.

Should there exist suitably structured optimal policies for the above individual location problems, then *location indices* may be developed which in turn will form the basis of good heuristics for the original problem in (1). To state the key requirement, write $\pi_l(\nu)$ for some stationary optimal policy for location l , achieving $C_l^{\text{opt}}(\nu)$. We also write $P\{\pi_l(\nu)\}$ for the *passive set* under $\pi_l(\nu)$, namely, the collection of inventory levels at location l for which $\pi_l(\nu)$ does not mandate a delivery.

DEFINITION 1. Location l is *indexable* if there exists a family of stationary optimal policies $\{\pi_l(\nu), \nu \in \mathbb{R}\}$ such that $P\{\pi_l(\nu)\}$ is increasing in ν . The corresponding *location l index* at inventory level J is defined by

$$\nu_l(J) = \inf\{\nu; \nu \in \mathbb{R} \text{ and } J \in P\{\pi_l(\nu)\}\}. \quad (7)$$

First, note that the requirement of indexability is simply that as the resource charge ν grows, then an optimal policy for location l will choose to deliver in fewer states. This seems reasonable. Second, the index $\nu_l(J)$ may be understood as a *fair charge* (per unit of resource) for a delivery to location l when its inventory level at the start of the day in which the delivery is made is J .

From Definition 1 and the earlier discussion surrounding (5) and (6), it follows that an optimal policy for the Lagrangian relaxation in (5) is structured as follows when all locations are indexable: at the start of each day, observe the inventory level at each location and mandate deliveries to all locations whose index (fair charge) is no less than ν (the prevailing charge). Before proceeding to describe how to use these indices to develop delivery heuristics for the original problem in (1), note that we can introduce a *null location* as one which when activated consumes some resource but which incurs no delivery or inventory costs. This null location represents a “no delivery” option. Any such location is trivially indexable with an index identically zero. This zero index value provides us with a suitable cutoff point when considering locations for delivery.

Heuristic Development

We follow Whittle (1988) in using the above index-based solution to the Lagrangian relaxation in (5) to argue for delivery heuristics which favour locations with large indices

for problems in which all locations are indexable. We shall call our standard proposal the *greedy index heuristic* (GI) which works as follows: at the start of each day, observe the inventory level (J_l) at each location (l) and compute the value of each location index ($\nu_l(J_l)$). Renumber the locations in decreasing order of their index values, such that

$$\nu_1(J_1) \geq \nu_2(J_2) \geq \dots \geq \nu_L(J_L). \quad (8)$$

Locations are considered for delivery only if they have a positive index. Hence, if $\nu_1(J_1) \leq 0$, no deliveries are made. Suppose that $\nu_1(J_1) > 0$. In this event, GI constructs a list of locations for delivery by first including location 1 and then considering the remaining locations with positive indices for inclusion in numerical order. When a location is considered, it is added to the list provided that the resulting collection of locations is feasible and not otherwise. A natural alternative, the *total index heuristic* (TI), chooses a collection of locations for delivery which is feasible and is such that the sum of the location indices is maximal. For the problems discussed in §5, GI and TI coincide.

We proceed to show in the next section that, subject to only very minor technical conditions, locations are indeed guaranteed indexable provided that the holding cost rate h is small enough and that the shortage cost σ exceeds the purchase price C .

3. Indexability Analysis

In this section, we focus on a single location and are thus able to drop the location identifier from the notation. Hence, the key inventory cost parameters are K , h , C , and σ , the distribution of daily demand is $\{p_j, j \geq 0\}$ with mean (rate) λ , the replenishment level is S , and the delivery time is τ . From §2, the single-location problem generated by the Lagrangian relaxation in (5) and fundamental to the study of indexability is as follows: determine a policy π for replenishing the location to minimise a combination of inventory costs incurred per unit of time ($C(\pi)$) and the delivery penalty rate ($\nu\tau I(\pi)$). We focus initially on the no holding costs case ($h = 0$) and make the details of the resulting single-location problem with multiplier ν explicit as follows:

(a) At the start of each day $n \in \mathbb{N}$, observe inventory level $i(n)$ and choose between action a (make a delivery) and action b (do not make a delivery). We introduce indicator I which takes the value one when action a is taken, and zero otherwise.

(b) If $i(n) = j \geq 0$ and a demand of k is experienced on day n , then the cost incurred on day n is

$$(K + \nu\tau)I + Ck$$

if $0 \leq k \leq j$ and is

$$(K + \nu\tau)I + Cj + \sigma(k - j)$$

if $j + 1 \leq k$. If $i(n) = j < 0$ (namely, that the location has already experienced j demands which it cannot meet)

and a demand of k is experienced on day n , then the cost incurred on day n is

$$(K + \nu\tau)I + \sigma k.$$

(c) If action a is taken on day n , then $i(n+1) = S$. If action b is taken on day n when $i(n) = j$, then $i(n+1) = j - k$, where k is the demand experienced on day n .

An optimal policy is a rule for taking actions which minimizes the average cost rate incurred over an infinite horizon. By standard theory (see, for example, Puterman 1994), we may restrict our search to *stationary policies*. Further, it is clear from the structure of (a)–(c) above that there will always be a single choice of optimal action for all nonpositive inventory levels. Indeed, the state space for the single-location problem can effectively be reduced to $\{0, 1, \dots, S\}$ in an obvious way. Hence, in our search for an optimal policy, we may restrict to maps from state space $\{0, 1, \dots, S\}$ to the action space $\{a, b\}$. However, we can say more. The following is a straightforward consequence of Theorem 8.11.3 of Puterman (1994):

LEMMA 1 (OPTIMALITY OF MONOTONE POLICIES). *If $h = 0$ and $\sigma > C$, then $\forall \nu \in \mathbb{R}$, $\exists J_\nu \in \{-\infty\} \cup \{0, 1, \dots, S\}$ such that the policy*

$$\pi(j) = a \iff j \leq J_\nu \quad (9)$$

is optimal for the single-location problem with multiplier ν .

Comment

Please note that the choice $J_\nu = -\infty$ in (9) corresponds to the policy which never makes a delivery.

In the following discussion, we restrict to the case $h = 0$, $\sigma > C$. In Lemma 1, in the event that there are several thresholds J_ν achieving the optimum for some ν , we write $\tilde{J}_\nu$ for the smallest. From Definition 1 in §2, we have *indexability* for the location if $\tilde{J}_\nu$ is decreasing in ν . To describe the resulting index, we introduce the function $\Psi: \mathbb{Z} \rightarrow \mathbb{R}$, given by

$$\Psi(j) = \begin{cases} (\sigma - C) \sum_{k=j+1}^{\infty} (k - j)p_k, & j \geq 0, \\ \Psi(0) - (\sigma - C)j, & j < 0. \end{cases} \quad (10)$$

It is easy to show that Ψ is positive, decreasing, and convex.

Consider now the monotone policy for the single-location problem with threshold $J \in \{0, 1, \dots, S\}$, namely, that which mandates a delivery (at the start of the day) as soon as the inventory level is at or below J . Plainly, the inventory-level process under policy J (as we shall call it) regenerates upon every entry into state S . We refer to the period between successive entries into S as a *single cycle* of the process. At present use, $\{\bar{P}_j(J), j \in (-\infty, J]\}$ for the probability distribution of the inventory level when a delivery is mandated

under policy J , with a negative inventory level representing a shortage. This may be thought of as the point of first entry into the set $(-\infty, J]$ for the inventory-level process when measured at the end of each day. The expected inventory/delivery costs incurred in a single cycle of the process under policy J may now be written as

$$\begin{aligned}\tilde{C}(J) &= \sum_{j=0}^J \bar{P}_j(J) \{C(S-j) + \phi(j)\} \\ &\quad + \sum_{j=1}^{\infty} \bar{P}_{-j}(J) \{CS + \phi(0) + \sigma j\} + K, \quad (11)\end{aligned}$$

where

$$\phi(j) = \sum_{k=j+1}^{\infty} \sigma(k-j)p_k + \sum_{k=0}^j Ck p_k + Cj \sum_{k=j+1}^{\infty} p_k, \quad j \in \mathbb{N}.$$

The quantity $\phi(j)$ may be understood as the expected inventory cost experienced on the day of delivery when the inventory level at the start of the day is j .

A formula for the expected cycle time may now be inferred from the martingale stopping theorem. If we write $i(n)$ for the inventory level at the start of day n , then under a policy of no replenishment, the process $\{i(n) + \lambda n, n \geq 0\}$ is plainly a martingale. Suppose now that $i(0) = S > J$, and write

$$n(J) = \min\{n \geq 0; i(n) \leq J\}.$$

We take $n(S) = 0$. By the martingale stopping theorem, we have

$$E\{i[n(J)]\} = S - \lambda E\{n(J)\}. \quad (12)$$

But we have that

$$E\{i[n(J)]\} = \sum_{j=0}^J j \bar{P}_j(J) + \sum_{j=1}^{\infty} (-j) \bar{P}_{-j}(J),$$

and hence from (12) that the mean length of a delivery cycle under policy J is given by

$$\begin{aligned}\tilde{T}(J) &= 1 + E\{n(J)\} \\ &= 1 + \lambda^{-1} \left\{ \sum_{j=0}^J (S-j) \bar{P}_j(J) + \sum_{j=1}^{\infty} (S+j) \bar{P}_{-j}(J) \right\}. \quad (13)\end{aligned}$$

We can now use (11) and (13) to write the total of the inventory costs and the penalty costs incurred under policy J per unit of time as

$$\gamma(J, \nu) = \{\tilde{C}(J) + \nu \tau\} \{\tilde{T}(J)\}^{-1} \quad (14)$$

$$\begin{aligned}&= \left(\sum_{j=0}^J \bar{P}_j(J) \{C(S-j) + \phi(j)\} \right. \\ &\quad \left. + \sum_{j=1}^{\infty} \bar{P}_{-j}(J) \{CS + \phi(0) + \sigma j\} + K + \nu \tau \right) \\ &\quad \cdot \left(1 + \lambda^{-1} \left\{ \sum_{j=0}^J (S-j) \bar{P}_j(J) + \sum_{j=1}^{\infty} (S+j) \bar{P}_{-j}(J) \right\} \right)^{-1} \\ &= C\lambda + \lambda \left[\frac{\sum_{j=0}^J \Psi(j) \bar{P}_j(J) + \sum_{j=1}^{\infty} \Psi(-j) \bar{P}_{-j}(J) + K + \nu \tau}{\lambda + \sum_{j=0}^J (S-j) \bar{P}_j(J) + \sum_{j=1}^{\infty} (S+j) \bar{P}_{-j}(J)} \right]. \quad (15)\end{aligned}$$

It is plain from the definitions of the quantities concerned that

$$\tilde{T}(J-1) \geq \tilde{T}(J), \quad 1 \leq J \leq S.$$

If this inequality is strict, then for J in the range $1 \leq J \leq S$ there is a unique ν -solution to the equation

$$\gamma(J, \nu) = \gamma(J-1, \nu),$$

which we shall call $\nu(J)$. From (14), if $\tilde{T}(J-1) > \tilde{T}(J)$, we have that

$$\begin{aligned}\nu(J) &= \tau^{-1} \{ \tilde{C}(J-1) \tilde{T}(J) - \tilde{C}(J) \tilde{T}(J-1) \} \\ &\quad \cdot \{ \tilde{T}(J-1) - \tilde{T}(J) \}^{-1}, \quad 1 \leq J \leq S.\end{aligned}$$

Additionally, we write $\nu(0)$ for the unique ν -solution to the equation

$$\gamma(0, \nu) = \sigma \lambda, \quad (16)$$

where trivially $\sigma \lambda$ is the cost rate incurred under the monotone policy with threshold $-\infty$ which never mandates a delivery.

THEOREM 1 (INDEXABILITY CRITERION). *If*

- (i) $\tilde{T}(J-1) > \tilde{T}(J)$, $1 \leq J \leq S$, and
- (ii) $\nu(J-1) > \nu(J)$, $1 \leq J \leq S$,

then the location is indexable with $\nu(J)$ the index for inventory level J .

PROOF. By Lemma 1, the search for optimal delivery policies may be restricted to the monotone class with thresholds chosen from the collection $\{-\infty\} \cup \{0, 1, \dots, S\}$. It is trivial from (16) that

$$\gamma(0, \nu) > \sigma \lambda \iff \nu > \nu(0). \quad (17)$$

Further, simple algebra serves to show that, under condition (i), for any J in the range $0 \leq J \leq S-1$,

$$\nu > \nu(J+1) \implies \gamma(J+1, \nu) > \gamma(J, \nu). \quad (18)$$

But by condition (ii) and (18), for any J in the range $0 \leq J \leq S-1$,

$$\begin{aligned}\nu > \nu(J+1) &\implies \nu > \nu(j) \\ &\implies \gamma(j, \nu) > \gamma(j-1, \nu), \\ &\hspace{15em} J+1 \leq j \leq S. \quad (19)\end{aligned}$$

We also have that, under condition (i), for any J in the range $1 \leq J \leq S$,

$$\nu < \nu(J) \implies \gamma(J, \nu) < \gamma(J-1, \nu). \quad (20)$$

But by condition (ii) and (20), for any J in the range $1 \leq J \leq S$,

$$\begin{aligned} \nu < \nu(J) &\implies \nu < \nu(j) \\ &\implies \gamma(j, \nu) < \gamma(j-1, \nu), \quad 1 \leq j \leq J. \end{aligned} \quad (21)$$

From (17), (19), and (21), it is trivial to infer that

$$\gamma(J, \nu) < \min \left[\left\{ \min_{0 \leq j \neq J \leq S} \gamma(j, \nu) \right\}, \sigma\lambda \right],$$

and hence that monotone policy J is strictly optimal for the single-location problem with penalty ν for the range $\nu(J+1) < \nu < \nu(J)$, $0 \leq J \leq S-1$. In this range of J , if $\nu = \nu(J+1)$, then both monotone policies J and $J+1$ are optimal. Further, if $\nu > \nu(0)$, then monotone policy $-\infty$ (never deliver) is strictly optimal. If $\nu = \nu(0)$, then both monotone policies $-\infty$ and zero are optimal. Finally, if $\nu < \nu(S)$, monotone policy S (deliver every day) is strictly optimal. If $\nu = \nu(S)$, then both monotone policies $S-1$ and S are optimal.

It follows from the above characterisation of optimal policies that the requirements of indexability in Definition 1 are met. Further, for all J in the range $0 \leq J \leq S$, the quantity $\nu(J)$ is indeed the infimum of all ν -values for which inventory level J is in the passive (do not deliver) set. It then follows from Definition 1 that $\nu(J)$ is the index for J . This concludes the proof. $\square$

To develop things further, suppose that the inventory level at the beginning of day 0 is S and use X_n for the demand experienced on day $n-1$, $n \geq 1$. Write

$$p^{(n)}(j) := P(X_1 + X_2 + \cdots + X_n = j), \quad j \in \mathbb{N}$$

for the n -fold convolution of the daily demand distribution. We now have, utilising the assumption that demands on different days are independent, that

$$\begin{aligned} \bar{P}_j(J) &= P(X_1 = S-j) + \sum_{n=1}^{\infty} \sum_{r=0}^{S-J-1} P(X_1 + \cdots + X_n = r) \\ &\quad \cdot P(X_{n+1} = S-j-r) \\ &= p_{S-j} + \sum_{n=1}^{\infty} \sum_{r=0}^{S-J-1} p^{(n)}(r) p_{S-j-r}, \quad j \in (-\infty, J]. \end{aligned} \quad (22)$$

It will simplify matters in what follows if we write

$$P_j(J) = \bar{P}_{J-j}(J), \quad j \in \mathbb{N}. \quad (23)$$

Hence, we have from (22) that

$$P_j(J) = p_{S-J+j} + \sum_{n=1}^{\infty} \sum_{r=0}^{S-J-1} p^{(n)}(r) p_{S-J+j-r}, \quad j \in \mathbb{N}. \quad (24)$$

We now state our main result.

THEOREM 2 (INDEXABILITY AND LOCATION INDEX). *If $h=0$, $\sigma > C$, and the demand distribution $\{p_j, j \geq 0\}$ is such that*

$$\sum_{n=1}^{\infty} p^{(n)}(S-J) > 0, \quad 0 \leq J \leq S, \quad (25)$$

then the location is indexable and the index at inventory level J is given by

$$\begin{aligned} \nu(J) &= -K(\tau)^{-1} + (\lambda\tau)^{-1} \left[\left\{ \sum_{j=1}^{\infty} p_j [\Psi(J-j) - \Psi(J)] \right\} \right. \\ &\quad \cdot \left. \left\{ \lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) \right\} \right] \\ &\quad - (\tau)^{-1} \sum_{j=0}^{\infty} \Psi(J-j)P_j(J), \quad 0 \leq J \leq S. \end{aligned} \quad (26)$$

To establish Theorem 2, we first observe from (15) and (23) that the cost rate $\gamma(J, \nu)$ may be re-expressed as

$$\gamma(J, \nu) = C\lambda + \lambda \left\{ \frac{\sum_{j=0}^{\infty} \Psi(J-j)P_j(J) + K + \nu\tau}{\lambda + S - J + \sum_{j=0}^{\infty} jP_j(J)} \right\}, \quad 0 \leq J \leq S, \quad (27)$$

with $P_j(J)$ given by (24). Note that it is trivial to establish from (24) that

$$\begin{aligned} P_{j-1}(J-1) - P_j(J) &= p_j \sum_{n=1}^{\infty} p^{(n)}(S-J), \\ j &\in \mathbb{Z}^+, \quad 0 \leq J \leq S, \end{aligned} \quad (28)$$

and it will be convenient to write

$$\sum_{n=1}^{\infty} p^{(n)}(S-J) = \delta(S-J), \quad 0 \leq J \leq S, \quad (29)$$

in what follows. Note from (25) and (29) that the condition in Theorem 2 may be expressed as

$$\delta(S-J) > 0, \quad 0 \leq J \leq S.$$

We easily deduce from (28) and (29) that, for any $0 \leq J \leq S$,

$$\sum_{j=0}^{\infty} jP_j(J-1) - \sum_{j=0}^{\infty} jP_j(J) = \lambda\delta(S-J) - 1,$$

and hence from (13) and (23) that

$$\begin{aligned} \lambda\tilde{T}(J) &= \lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) \\ &= \lambda + S - J + 1 + \sum_{j=0}^{\infty} jP_j(J-1) - \lambda\delta(S-J) \\ &= \lambda\tilde{T}(J-1) - \lambda\delta(S-J), \quad 1 \leq J \leq S. \end{aligned} \quad (30)$$

Hence, the requirement of Theorem 1(i) is met under the condition on the demand distribution $\{p_j, j \geq 0\}$ given in Theorem 2. We shall state as much in the following result.

LEMMA 2. If the demand distribution satisfies condition (25), then

$$\tilde{T}(J-1) > \tilde{T}(J), \quad 1 \leq J \leq S.$$

From Theorem 1, to establish Theorem 2 it remains to prove that

$$\nu(J-1) > \nu(J), \quad 1 \leq J \leq S,$$

where for the range $1 \leq J \leq S$, $\nu(J)$ is the ν -value for which

$$\gamma(J, \nu) = \gamma(J-1, \nu).$$

From (27)–(30), we infer that for any $1 \leq J \leq S$,

$$\begin{aligned} & \gamma(J-1, \nu) \\ &= C\lambda + \lambda \left\{ \frac{\sum_{j=1}^{\infty} \Psi(J-j)P_j(J) + \delta(S-J)\sum_{j=1}^{\infty} \Psi(J-j)p_j + K + \nu\tau}{\lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) + \lambda\delta(S-J)} \right\}. \end{aligned} \quad (31)$$

If we equate the expressions in (27) and (31) and solve for ν , then it is straightforward to show that for the range $1 \leq J \leq S$, we obtain the expression for $\nu(J)$ given in (26). It is also straightforward to show that $\nu(0)$ is given by the appropriate form of expression (26). If we now use (26), (28), and (30), we recover the expression

$$\begin{aligned} & \nu(J-1) \\ &= -K(\tau)^{-1} + (\lambda\tau)^{-1} \left[\left\{ \sum_{j=1}^{\infty} p_j [\Psi(J-1-j) - \Psi(J-1)] \right\} \right. \\ & \quad \times \left\{ \lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) + \lambda\delta(S-J) \right\} \\ & \quad \left. - (\tau)^{-1} \left\{ \sum_{j=1}^{\infty} \Psi(J-j)[P_j(J) + p_j\delta(S-J)] \right\} \right] \\ &= -K(\tau)^{-1} + (\lambda\tau)^{-1} \left[\left\{ \sum_{j=1}^{\infty} p_j [\Psi(J-1-j) - \Psi(J-1)] \right\} \right. \\ & \quad \times \left\{ \lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) \right\} \\ & \quad \left. - (\tau)^{-1} \sum_{j=0}^{\infty} \Psi(J-j)P_j(J) + (\tau)^{-1} \Psi(J)P_0(J) \right. \\ & \quad \left. + (\tau)^{-1} \delta(S-J) \left\{ \sum_{j=1}^{\infty} p_j [\Psi(J-1-j) - \Psi(J-1)] \right\} \right. \\ & \quad \left. - (\tau)^{-1} \delta(S-J) \sum_{j=1}^{\infty} p_j \Psi(J-j), \quad 1 \leq J \leq S. \right] \end{aligned} \quad (32)$$

However, we have from (24) that

$$P_0(J) = p_{S-J} + \sum_{n=1}^{\infty} \sum_{r=0}^{S-J-1} p^{(n)}(r) p_{S-J-r},$$

and hence that

$$\begin{aligned} & P_0(J) \{\delta(S-J)\}^{-1} + p_0 \\ &= \left\{ p_{S-J} + \sum_{n=1}^{\infty} \sum_{r=0}^{S-J} p^{(n)}(r) p_{S-J-r} \right\} \{\delta(S-J)\}^{-1} = 1. \end{aligned}$$

It then follows that

$$P_0(J) = (1 - p_0)\delta(S-J), \quad (33)$$

and so from (25), (26), (32), (33) and the strict convexity of Ψ over the range $[-1, \infty)$, we deduce that

$$\begin{aligned} & \nu(J-1) > -K(\tau)^{-1} + (\lambda\tau)^{-1} \left[\left\{ \sum_{j=1}^{\infty} p_j [\Psi(J-j) - \Psi(J)] \right\} \right. \\ & \quad \cdot \left\{ \lambda + S - J + \sum_{j=0}^{\infty} jP_j(J) \right\} \\ & \quad \left. - (\tau)^{-1} \sum_{j=0}^{\infty} \Psi(J-j)P_j(J) + (\tau)^{-1} \delta(S-J) \right. \\ & \quad \cdot \left[\sum_{j=1}^{\infty} p_j \{ [\Psi(J-1-j) - \Psi(J-1)] \right. \\ & \quad \left. \left. - [\Psi(J-j) - \Psi(J)] \} \right] \right] > \nu(J), \\ & \quad 1 \leq J \leq S. \end{aligned} \quad (34)$$

Theorem 2 now follows from Theorem 1, Lemma 2, and inequality (34).

Comments and Examples

1. Condition (25) is trivially satisfied for any demand distribution such that $\min(p_0, p_1) > 0$.

2. We specialise to the important case in which demand arises at the location according to a Poisson process with rate λ per day by taking

$$p_j = \frac{\lambda^j}{j!} e^{-\lambda}, \quad j \in \mathbb{N},$$

and

$$P_j(J) = \frac{\lambda^{S-J+j}}{(S-J+j)!} e^{-\lambda} \left[1 + \sum_{k=0}^{S-J-1} \binom{S-J+j}{k} \sum_{n=0}^{\infty} n^k e^{-n\lambda} \right], \quad j \in \mathbb{N}.$$

These may be substituted into (26) to obtain the index in this case.

3. One very simple case to analyse is that in which the daily demand has a geometric distribution with mean λ . Hence, we have

$$p_j = \lambda^j (1 + \lambda)^{-j-1}, \quad j \in \mathbb{N}. \quad (35)$$

It is clear from the memoryless property of the geometric distribution that the distribution of the inventory level upon the first entry in $(-\infty, J]$ from above must also have a geometric form, namely,

$$\bar{P}_j(J) = \lambda^{J-j}(1 + \lambda)^{j-J+1}, \quad j \in (-\infty, J].$$

These facts hugely simplify the calculations for this case. The index at inventory level J is given by

$$\nu(J) = -K(\tau)^{-1} + (\lambda\tau)^{-1}(\sigma - C)\{\lambda/1 + \lambda\}^{J+1} \cdot \{S\lambda + (S - J)(J + 1)\}. \quad (36)$$

4. From inequality (34), the index $\nu(J)$ is strictly decreasing in the inventory level J . This fact is fundamental to the application of Theorem 1. It is clear that for any reasonable approach to the modelling of holding costs, the decreasing nature of the index must remain the case if holding costs are introduced provided only that the holding cost parameter h is small enough. We consider now the important case of Poisson demand. For h small enough, we obtain a new index $\nu_h(J)$ by incorporating holding costs into the above calculations. The new index is as follows:

$$\nu_h(J) = \nu(J) + (\lambda\tau)^{-1}\{H(J - 1) - H(J)\} - (\tau)^{-1}H(J)\delta(S - J), \quad 0 \leq j \leq S, \quad (37)$$

where

$$H(J) = h \sum_{n=1}^S \sum_{j'=J+1}^S \left[\sum_{j=0}^J P(n, j', j) \left\{ \frac{(S+j')(n-1)}{2} + \frac{(j+j')}{2} + \Phi(j) \right\} + \sum_{j \leq -1} P(n, j', j) \left\{ \frac{(S+j')(n-1)}{2} + \frac{j'(j'+1)}{2(j'-j+1)} \right\} \right]. \quad (38)$$

In (38), the quantities $P(n, j', j)$ and $\Phi(j)$ are given by

$$P(n, j', j) = \frac{\lambda^{S-j} e^{-\lambda n} (n-1)^{S-j'}}{(S-j')!(j'-j)!}, \quad n \geq 1, S \geq j' \geq J+1, J \geq j,$$

and

$$\Phi(j) = \sum_{k=0}^j \left(j - \frac{k}{2} \right) \frac{\lambda^k}{k!} e^{-\lambda} + \sum_{k=j+1}^{\infty} \frac{j(j+1)}{2(k+1)} \frac{\lambda^k}{k!} e^{-\lambda}, \quad J \geq j \geq 1.$$

In practice, we need that the index function $\nu_h(J)$ be decreasing in J for the range of J for which it is positive. Recall that making no delivery is preferable to making one to a negative index location. We have found that this more limited requirement for the index function is satisfied for problems in which the holding cost parameter h takes practically realistic values.

Table 1. Values of the index $\nu_h(J)$ for a single location (details in text).

J	0	1	2	3	4
$\nu_h(J)$	398.62	398.62	398.62	398.62	398.61
J	5	6	7	8	9
$\nu_h(J)$	398.55	398.35	397.82	396.58	394.01
J	10	11	12	13	14
$\nu_h(J)$	389.25	381.24	368.85	351.05	327.11
J	15	16	17	18	19
$\nu_h(J)$	296.72	260.04	217.67	170.54	119.80
J	20	21	22	23	24
$\nu_h(J)$	66.65	12.31	-42.12	-95.61	-147.28

As an illustration of the above, find in Table 1 values of the index $\nu_h(J)$ over the range $0 \leq J \leq 24$ for a location facing Poisson demand with rate $\lambda = 15$ and for which $S = 90$. The inventory cost parameters are $\sigma = 20$, $C = 10$, and $h = 0.01$. Note that if the time unit is taken as one day, then a standard approach to the setting of h would be to take

$$h = \alpha C / 365$$

with α a suitably chosen (annualised) interest rate. The above choices would correspond to a value of α equal to 36.5% per annum and hence this choice of h is actually larger than may be required in practice. We take the delivery time τ to be one and the fixed delivery cost to be 500.

Note that it is clear from Table 1 that the quantity $\nu_h(J)$ comfortably meets the indexability requirement that it be decreasing in J over the range for which it is positive. The cut off for deliveries (i.e., the inventory level at which a delivery would certainly not be mandated) is 22, which is a little more than 1.8 standard deviations above the mean for a single day's demand.

5. Two issues arise when computing location indices for general demand distributions. First, note from (26) that computation of $\nu(J)$ requires the values of $P_j(J)$, $j \in \mathbb{N}$. In the case of deterministic (zero variance) demand, this is trivial. If the variance of daily demand is nonzero, then appeal to the central limit theorem and to (24) means that we have

$$P_j(J) = p_{S-J+j} + \sum_{n=1}^N \sum_{r=0}^{S-J-1} p^{(n)}(r) p_{S-J+j-r} + A_j(N),$$

where $A_j(N) \rightarrow 0$, $N \rightarrow \infty$ at a geometric rate. Hence, $P_j(J)$ may be well approximated by small partial sums. The same applies to the computation of the quantities $\tilde{T}(J)$ and $\delta(S - J)$.

Second, as in comment 4 above, holding costs need to be incorporated into the index. Write $H(S', J)$, $J + 1 \leq S' \leq S$, for the expected holding cost incurred during a single cycle of the process under policy J , with replenishment level S' and under an approximating assumption that each day's

demand occurs uniformly. Conditioning upon the first day's demand following replenishment, we have the recursion

$$H(S', J) = \left\{ \sum_{k=1}^{S'-J-1} p_k H(S' - k, J) + \sum_{k=0}^{S'} p_k h \left(S' - \frac{k}{2} \right) + \sum_{k=S'+1}^{\infty} p_k h \frac{(S')^2}{2k} \right\} (1 - p_0)^{-1}, \quad J+1 \leq S' \leq S,$$

and this enables the computation of $H(S, J)$ for any $J \in \{0, 1, \dots, S\}$. The appropriate form of index is then given by

$$\nu_h(J) = \nu(J) + \{\tau \delta(S - J)\}^{-1} \{H(S, J - 1) - H(S, J)\} - (\tau)^{-1} H(S, J).$$

The above comments regarding indexability for realistic values of h continue to apply.

4. An Approximate (Continuously Observed) Index for the Poisson Case

We now specialise to the important case of Poisson demand and develop a simple approximation to the index of §3 derived under an assumption that the inventory level at each location is monitored continuously. More precisely, the single-location problem (a)–(c) described in §3 is modified to permit continuous observation of the inventory level. It is transparent from the Markovian nature of the Poisson demand process that in the resulting single-location problem decisions (deliver or not) need only to be taken at demand epochs and so the problem becomes semi-Markovian in nature. Under monotone policies, a delivery is triggered as soon as the inventory hits the appropriate threshold. To produce a process that approximates (a)–(c) well, we suppose that each delivery arrives one time unit after the demand epoch which triggered it. This quasi-lead time is hereafter referred to as the *delivery day*. The resulting approximation will enable us to develop a simpler index than that of §3 and will allow us to incorporate holding costs directly into the calculations, thereby permitting a direct analysis of the range of parameter values h which yield (approximate) indexability.

We consider a location with inventory cost parameters K , h , C , and σ and Poisson demand with rate λ . The expected duration of a single delivery cycle under monotone policy with threshold $J \in \{0, 1, \dots, S\}$ now simplifies to

$$\bar{T}(J) = 1 + \lambda^{-1}(S - J) \quad (39)$$

because a delivery is now triggered as soon as inventory level J is reached. The expected total inventory cost incurred during a single cycle (including holding costs) may be written as

$$\bar{C}(J) = \sigma \left\{ \sum_{j=J+1}^{\infty} (j - J) \frac{\lambda^j}{j!} e^{-\lambda} \right\} + C \left\{ S - \sum_{j=0}^J (J - j) \frac{\lambda^j}{j!} e^{-\lambda} \right\} + h \sum_{j=J+1}^S \frac{j}{\lambda} + h(J) + K. \quad (40)$$

The first term on the right-hand side of (40) is the expected cost incurred through shortages experienced during the delivery day, the second term is the expected cost of the items delivered, and the third term is the expected holding cost incurred during that part of the cycle which precedes the demand epoch triggering the delivery. Further, $h(J)$ is the expected holding cost incurred during the delivery day and is given by

$$h(J) = h \sum_{j=0}^J \left(J - \frac{j}{2} \right) \frac{\lambda^j}{j!} e^{-\lambda} + h \sum_{j=J+1}^{\infty} \frac{J(J+1)}{2(j+1)} \frac{\lambda^j}{j!} e^{-\lambda}. \quad (41)$$

To obtain (41), we exploit standard properties of the Poisson process. As in §3, we write the total of the inventory costs and the penalty costs incurred under policy J per unit of time as

$$\bar{\gamma}(J, \nu) = \{\bar{C}(J) + \nu \tau\} \{\bar{T}(J)\}^{-1}. \quad (42)$$

Note from (39) that it is plain that $\bar{T}(J - 1) > \bar{T}(J)$, $1 \leq J \leq S$. For J in the range $1 \leq J \leq S$, we write $\bar{\nu}(J)$ for the unique ν -solution to the equation

$$\bar{\gamma}(J, \nu) = \bar{\gamma}(J - 1, \nu). \quad (43)$$

Additionally, we write $\bar{\nu}(0)$ for the unique ν -solution to the equation

$$\bar{\gamma}(0, \nu) = \sigma \lambda.$$

We shall have indexability for this approximate analysis if $\bar{\nu}(J - 1) > \bar{\nu}(J)$, $1 \leq J \leq S$. The key facts are summarised in the next result. In the statement of Theorem 3, $[u]$ denotes the integer part of u .

THEOREM 3 (INDEXABILITY AND LOCATION INDEX—APPROXIMATE ANALYSIS FOR THE POISSON CASE). *The quantity $\bar{\nu}(J)$ is given by*

$$\begin{aligned} \bar{\nu}(J) = & -K(\tau)^{-1} + (\sigma - C)(\tau)^{-1} \left\{ S - \sum_{i=0}^{J-1} (S - i + \lambda) \frac{\lambda^i}{i!} e^{-\lambda} \right\} \\ & - h(2\lambda\tau)^{-1} S(S + 1) + h(\lambda\tau)^{-1} \\ & \cdot \sum_{i=0}^J \left\{ \frac{J(J+1)}{2} - \frac{i(i-1)}{2} + (S - J)(J - i) \right\} \frac{\lambda^i}{i!} e^{-\lambda}, \\ & 0 \leq J \leq S, \end{aligned} \quad (44)$$

and is strictly decreasing over the range $0 \leq J \leq [\lambda + \alpha\sqrt{\lambda}]$ where $\alpha > 0$, provided that

$$\sigma - C > h(\lambda)^{-1} \{1 + \lambda + \alpha\sqrt{\lambda}\} \exp\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right). \quad (45)$$

PROOF. The form of $\bar{\nu}(J)$ given in (44) follows from the expressions in (39)–(43) by means of straightforward algebra. Write

$$\Delta(J) := \tau\{\bar{\nu}(J + 1) - \bar{\nu}(J)\}.$$

It is easy to show from (44) that

$$\Delta(J) = -(\sigma - C)(S - J + \lambda) \frac{\lambda^J}{J!} e^{-\lambda} + h \left\{ \frac{\lambda^J}{J!} e^{-\lambda} + (\lambda)^{-1} \sum_{i=0}^J (S - J + i) \frac{\lambda^i}{i!} e^{-\lambda} \right\}. \quad (46)$$

Suppose that $J = \alpha\lambda$, where $\alpha \leq 1$. It then follows that

$$\frac{\lambda^i}{i!} e^{-\lambda} \leq \frac{\lambda^J}{J!} e^{-\lambda}, \quad 0 \leq i \leq J, \quad (47)$$

and hence that

$$\sum_{i=0}^J \frac{\lambda^i}{i!} e^{-\lambda} \leq (J+1) \frac{\lambda^J}{J!} e^{-\lambda} \quad (48)$$

and

$$\sum_{i=0}^J i \frac{\lambda^i}{i!} e^{-\lambda} \leq \lambda J \frac{\lambda^J}{J!} e^{-\lambda}. \quad (49)$$

Using (47)–(49) within (46), we see that

$$\Delta(J) \leq (S - J + \lambda) \frac{\lambda^J}{J!} e^{-\lambda} \left\{ -(\sigma - C) + \frac{h}{\lambda} + h\alpha \right\}. \quad (50)$$

It now follows from (50) that $\bar{\nu}(J)$ must be strictly decreasing over the range $0 \leq J \leq [\lambda + 1]$ when

$$(\sigma - C) > h(\lambda)^{-1}(1 + \lambda). \quad (51)$$

We now suppose that $J = [\lambda + \alpha\sqrt{\lambda}]$, where $\alpha > 0$. It will slightly simplify matters (although does not impact the result) if we suppose that $\lambda \in \mathbb{Z}^+$, and hence that $J = \lambda + [\alpha\sqrt{\lambda}]$. We then have that

$$\begin{aligned} \frac{\lambda^J}{J!} e^{-\lambda} &\geq \lambda^{[\alpha\sqrt{\lambda}]} \left\{ \prod_{i=1}^{[\alpha\sqrt{\lambda}]} (\lambda + i) \right\}^{-1} \frac{\lambda^\lambda}{\lambda!} e^{-\lambda} \\ &\geq \lambda^{[\alpha\sqrt{\lambda}]} \left\{ \prod_{i=1}^{[\alpha\sqrt{\lambda}]} (\lambda + i) \right\}^{-1} \left\{ \sum_{i=0}^J \frac{\lambda^i}{i!} e^{-\lambda} \right\} (J+1)^{-1}. \end{aligned} \quad (52)$$

Note that inequality (52) uses the fact that when $\lambda \in \mathbb{Z}^+$, it is a mode of the Poisson(λ) distribution. If we now use the fact that geometric means are bounded above by arithmetic means, we have that

$$\begin{aligned} \left\{ \prod_{i=1}^{[\alpha\sqrt{\lambda}]} (\lambda + i) \right\} \lambda^{-[\alpha\sqrt{\lambda}]} &\leq \left\{ 1 + \frac{(1 + [\alpha\sqrt{\lambda}])}{2\lambda} \right\}^{[\alpha\sqrt{\lambda}]} \\ &\leq \exp\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right). \end{aligned} \quad (53)$$

Using (53) within (52), we deduce that

$$\frac{\lambda^J}{J!} e^{-\lambda} \geq \exp\left\{-\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right)\right\} \left\{ \sum_{i=0}^J \frac{\lambda^i}{i!} e^{-\lambda} \right\} (J+1)^{-1},$$

and hence that

$$\begin{aligned} \frac{\lambda^J}{J!} e^{-\lambda} + (\lambda)^{-1} \sum_{i=0}^J (S - J + i) \frac{\lambda^i}{i!} e^{-\lambda} \\ \leq (\lambda)^{-1} (S - J + \lambda)(J+1) \frac{\lambda^J}{J!} e^{-\lambda} \exp\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right). \end{aligned} \quad (54)$$

From (46) and (54), we infer that

$$\begin{aligned} \Delta(J) &\leq (S - J + \lambda) \frac{\lambda^J}{J!} e^{-\lambda} \\ &\quad \cdot \left\{ -(\sigma - C) + \frac{h}{\lambda} (1 + [\lambda + \alpha\sqrt{\lambda}]) \exp\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right) \right\}, \end{aligned}$$

and hence that $\bar{\nu}(J+1) < \bar{\nu}(J)$ when

$$\sigma - C > h(\lambda)^{-1} \{1 + \lambda + \alpha\sqrt{\lambda}\} \exp\left(\frac{\alpha^2}{2} + \frac{\alpha}{2\sqrt{\lambda}}\right).$$

The result now follows easily. $\square$

Comments

1. As noted in §3, we in practice need the index to be decreasing for the range of inventory levels for which it is positive. The condition given in (45) in practice guarantees this for realistic values of the inventory parameters σ , C , and h . As an illustration of the above, find in Table 2 values of the index $\bar{\nu}(J)$ for a single location with stochastic demand and cost characteristics identical to those for the case discussed at the end of §3 and in Table 1.

In fact, the decreasing nature of $\bar{\nu}(J)$ for the range $0 \leq J \leq 24$ is easily inferred from Theorem 3. Note that these indices derived from an approximation which supposes continuous observation of the location's inventory lie below the corresponding exact values recorded in Table 1 implying more conservative decision making (in the sense of mandating deliveries at higher inventory levels) in the latter. For example, the delivery cut off is now 15 which is just the mean daily demand. This conservatism is as to be expected because in the model of §3, any decision not to deliver remains in force until the next decision epoch when the position is reviewed again. Hence, account has to be taken of the demand likely to occur during the next period. This is not necessary when the inventory is observed continuously.

Table 2. Values of the index $\bar{\nu}(J)$ for a single location (details in text).

J	0	1	2	3	4
$\nu_h(J)$	397.27	397.27	397.27	397.23	397.05
J	5	6	7	8	9
$\nu_h(J)$	396.40	394.47	389.68	379.51	360.56
J	10	11	12	13	14
$\nu_h(J)$	329.55	283.38	221.08	144.04	56.10
J	15	16	17	18	19
$\nu_h(J)$	-37.09	-129.25	-214.68	-289.20	-350.59
J	20	21	22	23	24
$\nu_h(J)$	-398.48	-433.97	-459.00	-475.85	-486.68

2. It seems reasonable to conjecture that the indices in (26) and the zero holding cost versions of the indices in (44) will be increased (for each J) upon replacement of the daily demand distribution by another which is stochastically larger. This increase would reflect greater exposure to shortage costs from any positive inventory level. While it has not proved possible to establish this in general, it has been found to hold in simple cases. For example, the geometric demand distribution in (35) is stochastically increasing in λ , while the corresponding index in (36) has a λ -derivative equal to

$$(\sigma - C)(\tau)^{-1}\{\lambda^{J-1}(1 - \lambda)^{-J-2}J(J+1)(S - J + \lambda)\},$$

which is strictly positive for $J > 0$ and zero when $J = 0$. Further, the Poisson(λ) distribution is also stochastically increasing in λ . When $h = 0$, the λ -derivative of the approximating Poisson index in (44) is zero when $J = 0$ and is

$$(\sigma - C)(\tau)^{-1}\lambda e^{-\lambda}\left\{(S - J)\frac{\lambda^{J-1}}{(J-1)!} + J\frac{\lambda^J}{J!}\right\} > 0$$

otherwise. Because greater exposure to shortage costs will generally correspond to reduced exposure to holding costs, we would expect the above property to be compromised by increasing the value of h above zero.

5. Numerical Examples

We now report results obtained in the course of an extensive numerical investigation into the quality of heuristic policies for delivery determined by the location indices presented in the preceding two sections. In all cases reported, individual locations face Poisson demands, and all delivery times are equal to the unit of time (one day) as is the available resource M . Hence, a single delivery (at most) is mandated each day. In no case studied did any location fail to be indexable.

The following heuristics were evaluated in all problems studied:

Greedy Index (GI). At the beginning of each day, the inventory level at each location is observed, with J_l the level at location l , $1 \leq l \leq L$. Compute the index $\nu_{hl}(J_l)$ (see (37)) for each l . If $\max_l \nu_{hl}(J_l) > 0$, then a delivery is mandated to the location of maximal index. If $\max_l \nu_{hl}(J_l) \leq 0$, then no delivery is made.

Greedy Approximate Index (GAI). This is structured as GI, but now the location l index used is $\bar{\nu}_l(J_l)$, to be found at (44) above.

Days Remaining (DR). During each day a delivery is made to a location which has the smallest value of J_l/λ_l . The latter quantity may be understood as the mean number of days (in the absence of any deliveries) to a stockout at location l .

We initially studied a range of small (two and three location) problems for which application of conventional

Table 3. Details (demand rates and replenishment levels) for 10 locations.

l	1	2	3	4	5	6	7	8	9	10
λ_l	15	15	5	17.5	7.5	15	10	12.5	5	2.5
S_l	90	90	50	175	75	150	100	125	150	75

stochastic dynamic programming (DP) methods was possible, although expensive. Hence, within accuracy, the cost rates incurred under the application of GI, GAI, and DR could all be obtained by DP value iteration and compared directly to optimal. In over 150 problems studied, the cost rate of GI was never more than 0.88% above the optimal cost rate and in the majority of cases was indistinguishable from it. The equivalent worst case figure for GAI was 3.17%. In no single problem studied did GAI outperform GI. On occasion, the cost rate under DR exceeded the optimum by more than 40%.

The study of small problems was complemented by consideration of larger problems for which any application of conventional DP methodologies was impossible. For these problems, cost rate estimates were obtained by the application of Monte Carlo simulation. As an illustration, find in Table 3 details (values of λ_l , S_l) of 10 locations. In this early part of the discussion, we shall assume that all delivery and inventory costs for the locations are identical with $\sigma_l = 20$, $C_l = 10$, $h_l = 0.01$, $1 \leq l \leq 10$.

Note that this example is “in balance” in the sense that it is possible to design a deterministic delivery schedule which supplies each location at the “right” rate to achieve a balance between supply and demand—i.e., each location l receives a delivery once in every S_l/λ_l time periods. One such deterministic schedule operates on a 30-period cycle. Within each cycle, deliveries are made as described in Table 4.

In what follows, DET denotes repeated application of this schedule and this policy will be evaluated, along with GI, GAI, and DR. Note that the balance mentioned above should assist the cost performance of the policies DR and DET which mandate deliveries in every time period.

In Table 3, find estimates of the cost rates incurred under our heuristic policies for a range of assumed values of a common delivery cost K . Note that the standard errors of the cost rate estimates have not been included in Table 5. We performed sufficient runs of our simulation model to guarantee that these were all very small in comparison with the observed differences between the estimated cost rates under the respective policies.

Table 4. A deterministic delivery schedule for the 10 locations in Table 3.

Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Location	1	2	3	4	5	6	1	2	7	8	3	4	1	2	5
Period	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Location	6	7	8	1	2	3	4	5	6	1	2	7	8	9	10

Table 5. Cost rate estimates for four delivery heuristics applied to a balanced 10 location problem (details in text).

K	GI	GAI	DR	DET
500	1,594.13	1,608.68	1,601.78	1,626.18
550	1,635.90	1,649.00	1,651.78	1,676.18
600	1,677.87	1,690.29	1,701.78	1,726.18
650	1,718.76	1,728.30	1,751.78	1,776.18
700	1,761.50	1,769.30	1,801.78	1,826.18
750	1,799.85	1,808.13	1,851.78	1,876.18
800	1,835.77	1,841.16	1,901.78	1,926.18
850	1,869.90	1,874.98	1,951.78	1,976.18
900	1,903.06	1,907.55	2,001.78	2,026.18
950	1,923.63	1,927.91	2,051.78	2,076.18
1,000	1,943.84	1,947.75	2,101.78	2,126.18

From Table 5, observe that GI consistently outperforms GAI, although its cost rate advantage remains modest throughout. Policies DR and DET, assisted by the problem's balance, are indeed competitive for modest delivery costs but become less so as K increases.

We now render the problem unbalanced by increasing the replenishment levels from S_l (as given in Table 3) to $[S_l + 2\sqrt{S_l}]$. For example, the replenishment level at location 1 is now two standard deviations above the mean demand rate for six time periods. In comparison with the previous balanced example, fewer deliveries should be made. Heuristics GI and GAI have the capacity to make appropriate adjustments to delivery rates, while DR and DET do not. Hence, we should expect the cost rate advantage of the former over the latter to grow in comparison with those reported in Table 5. This is indeed the case. The new cost rate estimates may be found in Table 6. We see that the cost rates for DR and DET are modestly reduced in comparison with those reported in Table 5. This reflects the fact that the adverse effect on holding costs from increasing the replenishment levels is more than compensated by the benevolent effect on shortage penalties. The structure of GI and GAI means that they may similarly profit from increased replenishment levels while making fewer deliveries.

Table 6. Cost rate estimates for four delivery heuristics applied to an unbalanced 10 location problem (details in text).

K	GI	GAI	DR	DET
500	1,504.29	1,519.71	1,560.16	1,570.91
550	1,544.32	1,557.66	1,610.16	1,620.91
600	1,584.95	1,596.21	1,660.16	1,670.91
650	1,623.94	1,633.51	1,710.16	1,720.91
700	1,658.47	1,669.05	1,760.16	1,770.91
750	1,695.40	1,702.12	1,810.16	1,820.91
800	1,730.10	1,739.14	1,860.16	1,870.91
850	1,764.66	1,771.46	1,910.16	1,920.91
900	1,800.26	1,806.19	1,960.16	1,970.91
950	1,832.29	1,837.66	2,010.16	2,020.91
1,000	1,862.37	1,866.35	2,060.16	2,070.91

Table 7. A summary of cost rates incurred when four delivery heuristics are applied to 200 balanced 10 location problems (details in text).

h	GI	GAI	DR	DET
0.0025	1,746.16 (0.08)	1,754.93 (0.08)	1,801.81 (0.08)	1,826.32 (0.09)
0.005	1,747.57 (0.08)	1,756.17 (0.08)	1,803.29 (0.08)	1,827.79 (0.09)
0.0075	1,748.78 (0.08)	1,757.29 (0.08)	1,804.77 (0.08)	1,829.25 (0.09)
0.01	1,750.07 (0.08)	1,758.53 (0.08)	1,806.25 (0.08)	1,830.72 (0.09)

To make it clear that the broad conclusions of the above study are not dependent on the assumption of a common delivery cost, the balanced problem of Table 5 was studied with delivery costs for the locations now drawn independently from a uniform distribution on the interval [500, 1,000]. Other features remained as previously, save only the holding cost parameter h , which was set at one of the four values {0.0025, 0.005, 0.0075, 0.01}. Fifty problems (i.e., determined by the associated values of K_l , $1 \leq l \leq 10$) were randomly generated for each h -value. In all of the 200 problems generated in total, GI incurred the smallest cost rate, followed by GAI then DR, with DET always incurring the largest costs. In Table 7, find values of the average cost rates incurred under the four heuristics for the 50 problems randomly generated for each value of h . Because of the additional source of variability caused by the introduction of randomly chosen delivery costs, we include standard errors (in parentheses) for the cost rate estimates in the table.

In summary, we found throughout our investigation that GI performed strongly, whether in comparison to an optimal policy (in small problems) or competitor heuristics. While it consistently outperformed GAI, the differences in cost rates were usually small and the alternative index heuristic based on an approximation positing continuous observation of inventory levels is usually an acceptable alternative when demands are Poisson.

Comment

We reflect further on the examples discussed above with Poisson demand and with other details given in Table 3. We do this by generalising their salient features and then by considering an appropriate asymptotic regime. In this way, we are able to develop stronger insights regarding the relative status of the policies GI and DET.

Consider an L -location problem in which a single delivery (at most) is mandated each day. All locations experience Poisson demand. The problem is (for now assumed to be) “in balance” and there exists a T -period cycle during which each location l receives d_l deliveries, where d_l divides T and where $S_l d_l = \lambda_l T$, $1 \leq l \leq L$. In what follows, we say that a location is of type l if its demand rate

and cost characteristics are identical to those of location l , $1 \leq l \leq L$. We now scale this L -location problem up to one involving nLT locations (with nT independent locations of each type l) serviced by nT trucks, where $n \in \mathbb{N}$. It is easy to show that there is a deterministic replenishment policy (DET) in which each location of type l is replenished at intervals all exactly equal to $t_l = T/d_l$. This is accomplished by always assigning any truck numbered $rT + s$, $0 \leq r \leq n-1$, to a location which matches the type in position s of the T -period cycle above, $1 \leq s \leq T$. It is straightforward that the overall cost rate incurred under DET may be written as

$$nT \sum_{l=1}^L (K_l + C_l E[\min\{S_l, X_l(t_l)\}] + \sigma_l E[\{X_l(t_l) - S_l\}^+] + \bar{H}_l)(t_l)^{-1}, \quad (55)$$

where $u^+ = \max(u, 0)$, $X_l(t_l)$ is the demand at location l over t_l days, and where

$$\bar{H}_l = h_l \sum_{j=0}^{S_l} \left(S_l - \frac{j}{2}\right) \frac{(\lambda_l t_l)^j}{j!} e^{-\lambda_l t_l} + h_l \sum_{j=S_l+1}^{\infty} \frac{S_l(S_l+1)}{2(j+1)} \frac{(\lambda_l t_l)^j}{j!} e^{-\lambda_l t_l}$$

is the expected holding cost between successive deliveries at each type l location, $1 \leq l \leq L$. Now modify DET by imposing the requirement that each scheduled delivery to locations of type l is only made if the inventory level on the day concerned is J_l or less, where $0 \leq J_l \leq S_l$, $1 \leq l \leq L$. We write the corresponding cost rate as

$$nT \sum_{l=1}^L C_l(J_l, t_l),$$

which is equal to (55) when $J_l = S_l$, $1 \leq l \leq L$. If we write J_l^* , $1 \leq l \leq L$, for the cost minimising threshold values, then plainly

$$100 \left[\left\{ \sum_{l=1}^L C_l(S_l, t_l) \right\} \left\{ \sum_{l=1}^L C_l(J_l^*, t_l) \right\}^{-1} - 1 \right] \quad (56)$$

is the percentage cost excess of DET over its best modification based on thresholds and is a lower bound on the cost excess of DET over an optimal delivery policy for the nLT locations. The cost excess in (56) will become more serious as the delivery costs K_l increase and as we render the problem unbalanced by allowing the replenishment levels S_l to increase. Consideration of the asymptotic regime determined by the limit $n \rightarrow \infty$ will leave the cost excess in (56) unchanged.

Now formally add to the nLT locations above nT null locations of the kind described following Definition 1 above. This gives a “no delivery” option to each truck every day. The problem of optimally serving the nLT locations

with nT trucks, each of which makes (at most) one delivery per day, may be formulated as a *restless bandit* with $nT(L+1)$ projects (nT of them null), nT of which are made active at each period. The asymptotic regime determined by the limit $n \rightarrow \infty$ in which the proportion of projects activated remains constant is precisely that considered for restless bandits by Weber and Weiss (1990), who demonstrated that, subject to mild conditions, the index policy (here the greedy index heuristic) is asymptotically optimal.

Acknowledgments

The authors acknowledge the helpful comments of the referees and the support of the Engineering and Physical Sciences Research Council for this work through the award of grants GR/T08562/01 and GR/S45188/01.

References

- Ansell, P. S., K. D. Glazebrook, J. Niño-Mora, M. O’Keeffe. 2003. Whittle’s index policy for a multi-class queueing system with convex holding costs. *Math. Methods Oper. Res.* **57** 21–39.
- Barnes-Schuster, D., Y. Bassok. 1997. Direct shipping and the dynamic single-depot/multi-retailer system. *Eur. J. Oper. Res.* **101** 509–518.
- Federgruen, A., D. Simchi-Levi. 1995. Analysis of vehicle routing and inventory-routing problems. M. O. Ball, T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. *Network Routing. Handbooks in Operations Research and Management Science*, Vol. 8, Chapter 4. North-Holland, Amsterdam, 297–373.
- Gittins, J. C. 1979. Bandit processes and dynamic allocation indices (with discussion). *J. Roy. Statist. Soc.* **B41** 148–177.
- Gittins, J. C. 1989. *Multi-Armed Bandit Allocation Indices*. Wiley, Chichester, UK.
- Glazebrook, K. D., C. Kirkbride. 2007. Dynamic routing to heterogeneous collections of unreliable servers. *Queueing Systems* **55** 9–25.
- Glazebrook, K. D., C. Kirkbride, D. Ruiz-Hernandez. 2006. Spinning plates and squad systems—Policies for bi-directional restless bandits. *Adv. Appl. Probab.* **38** 95–115.
- Glazebrook, K. D., H. M. Mitchell, P. S. Ansell. 2005. Index policies for the maintenance of a collection of machines by a set of repairmen. *Eur. J. Oper. Res.* **165** 267–284.
- Glazebrook, K. D., J. Niño-Mora, P. S. Ansell. 2002. Index policies for a class of discounted restless bandits. *Adv. Appl. Probab.* **34** 754–774.
- Kleywegt, A. J., V. S. Nori, M. W. P. Savelsbergh. 2002. The stochastic inventory routing problem with direct deliveries. *Transportation Sci.* **36** 94–118.
- Mahajan, A., D. Teneketzis. 2008. Multi-armed bandit problems. A. Hero, D. Castanon, D. Cochran, K. Kastella, eds. *Foundations and Applications of Sensor Management*, Springer, 121–151.
- Niño-Mora, J. 2001. Restless bandits, partial conservation laws and indexability. *Adv. Appl. Probab.* **33** 76–98.
- Niño-Mora, J. 2002. Dynamic allocation indices for restless projects and queueing admission control: A polyhedral approach. *Math. Programming* **93** 361–413.
- Papadimitriou, C. H., J. N. Tsitsiklis. 1999. The complexity of optimal queueing network control. *Math. Oper. Res.* **24** 293–305.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York.
- Weber, R. R., G. Weiss. 1990. On an index policy for restless bandits. *J. Appl. Probab.* **27** 637–648.
- Weber, R. R., G. Weiss. 1991. Addendum to “On an index policy for restless bandits.” *Adv. Appl. Probab.* **23** 429–430.
- Whittle, P. 1988. Restless bandits: Activity allocation in a changing world. J. Gani, ed. *A Celebration of Applied Probability (J. Appl. Probab. Spec. 25A)*, Applied Probab. Trust, Sheffield, UK, 287–298.